

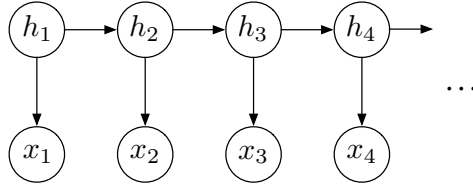
# Spectral Learning of Hidden Markov Model Predictive State Representations

Matthew James Johnson, MIT  
Revised March 2, 2013

This document is my summary explanation of the algorithm in “A Spectral Algorithm for Learning Hidden Markov Models” (COLT 2009), though there may be some slight notational inconsistencies with the original paper. The exposition and the math here are quite different, so if you don’t like this explanation, try the original paper!

The idea is to maintain output predictions in a recursive inference algorithm, instead of the usual method of maintaining hidden state predictions, and to represent the HMM only in terms of the maps necessary to update output predictions given new data. This approach limits the inference computations the algorithm can perform (it can’t answer any queries about the hidden states since it doesn’t explicitly deal with them at all), but it also reduces the complexity of the model parameters that are learned and thus makes learning easier. The learning algorithm uses an SVD and matrix operations, so it avoids the local-optima problems of EM or any other algorithms based on maximizing data likelihood over the usual HMM parameterization. The COLT paper includes error bounds and analysis.

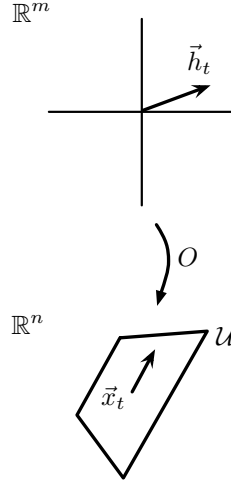
**Notation.** For a vector  $v$  in a subspace  $\mathcal{V} \subseteq \mathbb{R}^k$  and a matrix  $C$  with linearly independent columns and  $\text{range}(C) \supseteq \mathcal{V}$  I will use  $[v]^C$  to denote the coordinate vector of  $v$  relative to the ordered basis given by the columns of  $C$ , and  $[v]$  or simply  $v$  to denote the coordinate vector of  $v$  relative to the standard basis of  $\mathbb{R}^k$ . Similarly, for a linear map  $\mathcal{A} : \mathcal{V} \rightarrow \mathcal{V}$  I will use  $[\mathcal{A}]_C^C$  to denote the matrix of  $\mathcal{A}$  relative to domain and codomain bases given by the columns of  $C$ , and  $[\mathcal{A}]$  to indicate its matrix relative to standard bases. For a matrix  $A$  I will also use  $[A]_{ij}$  to denote the  $(i, j)$ th entry.



**Definition 1** (Hidden Markov Model). A time-homogeneous, discrete Hidden Markov Model (HMM) is a probability distribution on random variables  $\{(x_t, h_t)\}_{t \in \mathbb{N}}$  satisfying the conditional independences implied by the graphical model, where  $\text{range}(h_t) = [m] := \{1, 2, \dots, m\}$  and  $\text{range}(x_t) = [n]$  where  $n \geq m$ . The standard parameterization is the triple  $(T, O, \pi)$ , where

$$\begin{aligned} T &\in \mathbb{R}^{m \times m}, & [T]_{ij} &= \Pr[h_t = i | h_{t-1} = j] \\ O &\in \mathbb{R}^{n \times m}, & [O]_{ij} &= \Pr[x_t = i | h_t = j] \\ \pi &\in \mathbb{R}^m, & [\pi]_j &= \Pr[h_1 = j]. \end{aligned}$$

We will assume  $T$  and  $O$  to have full column rank and  $[\pi]_j > 0 \forall j \in [m]$ .



**Figure 1:** We can view  $O$  as (the matrix of) a map from hidden state beliefs to output predictions (with respect to the standard bases). Not shown is the fact that both  $\vec{h}_t$  and  $\vec{x}_t$  lie in the simplices of  $\mathbb{R}^m$  and  $\mathbb{R}^n$ , respectively, and that  $O$  maps the simplex in  $\mathbb{R}^m$  to (a subset of) the simplex in  $\mathbb{R}^n$ .

**Definition 2** (Observation Prediction). *An observation prediction for any time  $t$  is a vector  $\vec{x}_t \in \mathbb{R}^n$  defined in the standard basis by*

$$[\vec{x}_t]_i := \Pr[x_t = i | x_{1:t-1} = \bar{x}_{1:t-1}] \quad (1)$$

for some fixed (implicit) sequence  $\bar{x}_{1:t-1}$ .

**Claim 1.** *Every observation prediction  $\vec{x}_t$  lies in a subspace  $\mathcal{U} := \text{range}(O) \subseteq \mathbb{R}^n$  with  $\dim(\mathcal{U}) = m$ .*

*Proof.* By the conditional independences of the HMM, for any  $t$  we have

$$\begin{aligned} \Pr[x_t = i | x_{1:t-1} = \bar{x}_{1:t-1}] &= \sum_{j \in [m]} \Pr[x_t = i | h_t = j] \\ &\quad \cdot \Pr[h_t = j | x_{1:t-1} = \bar{x}_{1:t-1}] \end{aligned} \quad (2)$$

so therefore we can write

$$\vec{x}_t = O\vec{h}_t, \quad [\vec{h}_t]_j := \Pr[h_t = j | x_{1:t-1} = \bar{x}_{1:t-1}]. \quad (3)$$

Equivalently, we can say  $[\vec{x}_t]^O = \vec{h}_t$ . See Figure 1.  $\square$

**Claim 2.** *The observation prediction subspace  $\mathcal{U}$  satisfies  $\mathcal{U} = \text{range}(P_{2,1})$ , where  $[P_{2,1}]_{ij} := \Pr[x_2 = i, x_1 = j]$*

*Proof.* We can write the joint distribution over  $(x_1, x_2)$  as

$$\Pr[x_2 = i, x_1 = j] = \sum_{\bar{h}_1} \sum_{\bar{h}_2} \Pr[x_2 = i, x_1 = j, h_1 = \bar{h}_1, h_2 = \bar{h}_2] \quad (4)$$

$$\begin{aligned} &= \sum_{\bar{h}_1} \Pr[h_1 = \bar{h}_1] \Pr[x_1 = j | h_1 = \bar{h}_1] \\ &\quad \cdot \sum_{\bar{h}_2} \Pr[h_2 = \bar{h}_2 | h_1 = \bar{h}_1] \Pr[x_2 = i | h_2 = \bar{h}_2] \end{aligned} \quad (5)$$

and we can write that sum as  $P_{2,1} = OT \text{diag}^*(\pi)O^\top$ , where  $\text{diag}^*(\cdot)$  maps a vector to a diagonal matrix in the usual way. By our rank and positivity assumptions, we see that  $P_{2,1}$  satisfies  $\text{range}(P_{2,1}) = \text{range}(O) = \mathcal{U}$ .  $\square$

We can directly estimate  $P_{2,1}$  with empirical statistics, and as a consequence of Claim 2 we can then get a basis for  $\mathcal{U}$  by using an SVD of  $P_{2,1}$ . Note that we aren't getting an estimate of  $O$  this way, but just its column space.

**Claim 3.** *Given an observation  $x_t = \bar{x}_t$ , there is a linear map  $\mathcal{B}_{\bar{x}_t} : \mathcal{U} \rightarrow \mathcal{U}$  such that*

$$\mathcal{B}_{\bar{x}_t}(\vec{x}_t) = \alpha \vec{x}_{t+1} \quad (6)$$

for some  $\alpha = \alpha(\bar{x}_t, \vec{x}_t)$ , a scalar normalization factor chosen to ensure  $1^\top \vec{x}_{t+1} = 1$ .

*Proof.* Following the usual recursive update for HMM forward messages, we have

$$\begin{aligned} \Pr[h_{t+1} = i, x_{1:t} = \bar{x}_{1:t}] &= \sum_{j \in [m]} \underbrace{\Pr[h_{t+1} = i | h_t = j]}_{[T]_{ij}} \cdot \underbrace{\Pr[x_t = \bar{x}_t | h_t = j]}_{[O]_{\bar{x}_t j}} \\ &\quad \cdot \underbrace{\Pr[h_t = j, x_{1:t-1} = \bar{x}_{1:t-1}]}_{\propto [\bar{x}_t]_j^O}. \end{aligned} \quad (7)$$

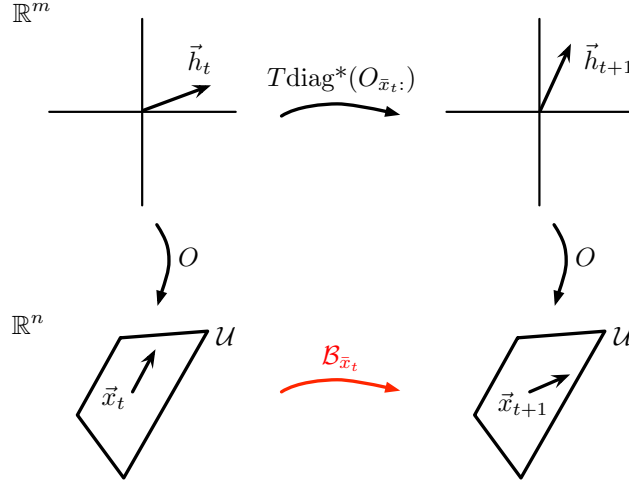
Therefore we can write the map  $\mathcal{B}_{\bar{x}_t}$  as an  $m \times m$  matrix relative to the basis for  $\mathcal{U}$  given by the columns of  $O$ :

$$[\mathcal{B}_{\bar{x}_t}]_O^O = T \text{diag}^*(O_{\bar{x}_t:}) \quad (8)$$

where  $O_{k:}$  denotes the vector formed by the  $k$ th row of  $O$ . We require  $1^\top \vec{x}_{t+1} = 1$ , so we have  $\alpha = O_{\bar{x}_t:}^\top \vec{x}_t$ .

See Figure 2.  $\square$

**Note.** Renormalization always works because  $T \text{diag}^*(O_{\bar{x}_t:})$  preserves the non-negative orthant of  $\mathbb{R}^m$ , and  $O$  maps the simplex in  $\mathbb{R}^m$  to (a subset of) the simplex in  $\mathbb{R}^n$ . The orthant preservation properties of these maps are an immediate consequence of the fact that the matrices (with respect to standard bases) are entry-wise non-negative. In fact, instead of tracking vectors, we should be tracking rays in the non-negative orthant.



**Figure 2:** The matrix  $T \text{diag}^*(O_{\bar{x}_t})$  is the matrix (relative to standard bases) of a linear map that updates hidden state beliefs given an observation  $\bar{x}_t$ , up to renormalization which the figure does not show. The linear map  $\mathcal{B}_{\bar{x}_t}$  is the update for output predictions.

For any new observation  $x$ , the map  $\mathcal{B}_x$  implements the “belief update” on output predictions, up to a normalization factor which we can compute on the fly via  $1^\top \vec{x}_t = 1$ . Note that we can also write  $\mathcal{B}_x$  as an  $n \times n$  matrix relative to the standard basis of  $\mathbb{R}^n$ :

$$[\mathcal{B}_x] = OT \text{diag}^*(O_{x\cdot})O^\dagger \quad (9)$$

where  $O^\dagger := (O^\top O)^{-1}O^\top$  is the pseudoinverse of  $O$ . Recall  $\vec{x}_t = O\vec{h}_t$  and hence  $\vec{h}_t = O^\dagger \vec{x}_t$ .

We would like to write  $\mathcal{B}_x$  as a matrix without reference to the standard HMM parameters  $(T, O, \pi)$ , since we want to avoid learning them at all.

**Claim 4.** *Let  $U \in \mathbb{R}^{n \times m}$  be a matrix whose columns form an orthonormal basis for  $\mathcal{U}$ . We can write  $\mathcal{B}_x$  as a matrix relative to the standard basis of  $\mathbb{R}^n$  as*

$$[\mathcal{B}_x] = P_{3,x,1}U(U^\top P_{2,1}U)^{-1}U^\top \quad (10)$$

where

$$[P_{3,x,1}]_{ij} := \Pr[x_3 = i, x_2 = x, x_1 = j] \forall x \in [n]. \quad (11)$$

*Proof.* We can express the matrix  $P_{3,x,1}$  in a form similar to that for the matrix  $[\mathcal{B}_x]$  in Equation (9):

$$P_{3,x,1} = \underbrace{OT \text{diag}^*(O_{x\cdot})O^\dagger}_{[\mathcal{B}_x]} P_{2,1}. \quad (12)$$

Intuitively, we want to remove the  $P_{2,1}$  on the right, since that would give us  $[\mathcal{B}_x]$  in terms of quantities we can readily estimate, but we cannot form the inverse of  $P_{2,1}$  because it is  $n \times n$  and has rank  $m \leq n$ . However,  $P_{2,1}$  has row and column space  $\mathcal{U}$  (intuitively, its restriction to  $\mathcal{U}$  is invertible), thus we can substitute

$$P_{2,1} = U(U^\top P_{2,1} U)U^\top \quad (13)$$

to get

$$P_{3,x,1} = OT \text{diag}^*(O_{x:})O^\dagger U(U^\top P_{2,1} U)U^\top \quad (14)$$

and hence

$$P_{3,x,1}U(U^\top P_{2,1} U)^{-1}U^\top = OT \text{diag}^*(O_{x:})O^\dagger = [\mathcal{B}_x]. \quad (15)$$

□

Because we can estimate each  $P_{3,x,1}$  as well as  $P_{2,1}$  from data by empirical statistics, and we can obtain a  $U$  using an SVD, we can now estimate a representation of  $\mathcal{B}_x$  from data using the expression in Claim 4. We can also directly estimate  $P_1$  from empirical statistics, where  $[P_1]_i := \Pr[x_1 = i]$ , and hence we can use these estimated quantities to recursively compute  $\Pr[x_t | x_{1:t-1}]$  and  $\Pr[x_{1:t-1}]$  given observations up to and including time  $t - 1$ .

Since  $\dim(U) = m \leq n$ , we can use the columns of  $U$  as our basis for  $\mathcal{U}$  to get a more economical coordinate representation of  $\vec{x}_t$  and  $\mathcal{B}_x$  than in the standard basis of  $\mathbb{R}^n$ :

**Definition 3** (HMM PSR Representation). *For any fixed  $U \in \mathbb{R}^{n \times m}$  with  $\text{range}(U) = \mathcal{U}$  and  $U^\top U = I_{m \times m}$ , we define the belief vector at time  $t$  by*

$$\vec{b}_t := [\vec{x}_t]^U = U^\top \vec{x}_t \in \mathbb{R}^m \quad (16)$$

and in particular for  $t = 1$  we have

$$\vec{b}_1 = U^\top O\pi. \quad (17)$$

For each possible observation  $x \in [n]$ , we define the matrix  $B_x \in \mathbb{R}^{m \times m}$  by

$$B_x := [\mathcal{B}_x]_U^U = (U^\top O)T \text{diag}^*(O_{x:})(U^\top O)^{-1}. \quad (18)$$

Finally, for normalization purposes, it is convenient to maintain the appropriate mapping of the ones (co-)vector, noting  $1^\top O \vec{h} = 1$  if and only if  $1^\top \vec{h} = 1$  because  $1^\top O = 1^\top$ :

$$\vec{b}_\infty := [1]^U = U^\top 1 \in \mathbb{R}^m. \quad (19)$$

The box below summarizes the method for learning an HMM PSR representation from data and how to use an HMM PSR representation to perform some recursive inference computations.

**Learning**

$$\begin{aligned}\hat{U} &= \text{ThinSVD}(\hat{P}_{2,1}) \\ \hat{b}_1 &= \hat{U}^\top \hat{P}_1 \\ \hat{B}_x &= \hat{U}^\top \hat{P}_{3,x,1} (\hat{U}^\top \hat{P}_{2,1})^\dagger \quad \forall x \in [n] \\ \hat{b}_\infty &= \hat{U}^\top \mathbf{1}\end{aligned}$$

**Inference**

$$\begin{aligned}\Pr[x_{1:t}] &= \vec{b}_\infty^\top B_{x_{t:1}} \vec{b}_1 && \text{sequence probability} \\ \Pr[x_t | x_{1:t-1}] &= \vec{b}_\infty^\top B_{x_t} \vec{b}_t && \text{prediction} \\ \vec{b}_{t+1} &= \frac{B_{x_t} \vec{b}_t}{\vec{b}_\infty^\top B_{x_t} \vec{b}_t} && \text{recursive update}\end{aligned}$$